arXiv:1510.02065v1 [cs.DC] 7 Oct 2015

Solving the Quadratic Assignment Problem on heterogeneous environment (CPUs and GPUs) with the application of Level 2 Reformulation and Linearization Technique

Alexandre Domingues Gonçalves
Campus São Gonçalo - Instituto Federal do Rio de Janeiro - IFRJ
Rua José Augusto Pereira dos Santos s/n - Neves - São Gonçalo - RJ
alexandre.domingues@ifrj.edu.br

Artur Alves Pessoa
Engenharia de Produção - Universidade Federal Fluminense - UFF
Rua Passo da Pátria 156 - Bloco E - 4º andar - São Domingos - Niterói - RJ
artur@producao.uff.br

Lúcia Maria de Assumpção Drummond
Instituto de Computação - Universidade Federal Fluminense - UFF
Rua Passo da Pátria 156 - Instituto de computação - sala 525 - São Domingos - Niterói - RJ
lucia@ic.uff.br

Cristiana Bentes
Departamento de Geomática - Universidade Estadual do Rio de Janeiro - UERJ
Rua São Francisco Xavier 524 - 5o Andar - Bloco D - Maracanã - Rio de Janeiro - RJ
cris@eng.uerj.com

Ricardo Farias
COPPE Sistemas - Universidade Federal do Rio de Janeiro - UFRJ
Centro de Tecnologia - Bloco H - Sala 319 - Ilha do Fundão - Rio de Janeiro - RJ
rfarias@cos.ufrj.br

ABSTRACT

The Quadratic Assignment Problem, QAP, is a classic combinatorial optimization problem, classified as NP-hard and widely studied. This problem consists in assigning N facilities to N locations obeying the relation of 1 to 1, aiming to minimize costs of the displacement between the facilities. The application of Reformulation and Linearization Technique, RLT, to the QAP leads to a tight linear relaxation but large and difficult to solve. Previous works based on level 3 RLT needed about 700GB of working memory to process one large instances (N = 30 facilities). We present a modified version of the algorithm proposed by Adams et al. which executes on heterogeneous systems (CPUs and GPUs), based on level 2 RLT. For some instances, our algorithm is up to 140 times faster and occupy 97% less memory than the level 3 RLT version. The proposed algorithm was able to solve by first time two instances: tai35b and tai40b.

KEYWORDS. QAP. RLT. GPU.

Main Area: Combinatorial Optimization

1. Introduction

The Quadratic Assignment Problem, or QAP, is one of the hardest and studied combinatorial optimization problems of literature. Consists to find an allocation of N facilities to N locations obeying the ratio of 1 to 1 and aiming to minimize the cost obtained by the sum of flow-distance products. Its formulation was initially presented by Koopmans e Beckmann (1957) and has practical applications in object allocation into departments, electronic circuit boards design, layouts problems, construction planning, and others. For a detailed study on the QAP we suggest the following references: Pardalos *et al*. (1994), Padberg e Rijal (1996), Burkard *et al*. (1998) and Cela (1997).

Exact methods to solve the QAP require high computational power, for example, in Hahn *et al*. (2013), one instance with 30 facilities (as nug30) spent 8 days in a supercomputer composed by 32 nodes of 2.2GHz Intel Xeon CPU and 700GB of main memory. The best known strategy to achieve the exact solution of QAP is using branch-and-bound algorithms where the main problem is divided into smaller sub-problems in trying to solve them exactly.

The lower bound are essential components for branch-and-bound procedures because they allow discards a large number of alternatives in the search for the optimal solution. The limits reached by Burer e Vandenbussche (2006), Adams *et al*. (2007) and Hahn *et al*. (2012) stand out as the most tighter to the QAP. The proposed by Burer e Vandenbussche (2006) consists of relaxations lift-and-project for binary integer problems, Adams *et al*. (2007) presents a dual ascent algorithm based on level 2 reformulation and linearization technique and Zhu *et al*. (2012) that also present a dual ascent algorithm but based on level 3 reformulation and linearization technique.

The reformulation and linearization technique, or RLT, was initially developed by Adams e Sherali (1990), Sherali e Adams (1999), in order to generate linear relaxations with tight lower bounds for a class of mixed integer programming 0-1 problems. In the literature, we find applied to the QAP the level 1 reformulation and linearization technique, or RLT1, in Hahn e Grant (1998), the level 2 reformulation and linearization technique, or RLT2, in Adams *et al*. (2007) and the level 3 reformulation and linearization technique, or RLT3 in Hahn *et al*. (2012). Parallel versions of RLT3 are proposed in Hahn *et al*. (2013) and Gonçalves *et al*. (2013) for a cluster with shared memory and a distributed environment, respectively. The mentioned studies have shown that higher the level of RLT used, tighter lower bound is obtained, however, the working memory (RAM) required to execute the algorithm increases exponentially as grows the level of RLT.

We didn't find studies in the literature addressing exact methods for the QAP solution using a computational structure composed of CPUs and GPUs. We found heuristics, such as the work of Tsutsui e Fujimoto (2011) based on ant colony optimization algorithm combined with Tabu search and another using genetic algorithms in Tsutsui e Fujimoto (2009).

Our proposal consists of an branch-and-bound algorithm that uses to calculate the lower bound a modified version of the dual ascent algorithm of Adams *et al*. (2007) to be executed on a heterogeneous environment (CPUs and GPUs). The algorithm takes advantage of the good bounds reached by RLT2 relaxation and the computational power of GPUs with low use of working memory when compared to RLT3 dual ascent algorithm. The RLT3 algorithm reaches tighter bounds than the previous but it requires a prohibitive memory for commercial computers.

This work is organized as follow: In the section 2 we show the QAP formulation. In the Section 3 we describe the RLT method and its application to QAP. In the Section 4 we present the dual ascent procedure based in RLT2 and its operations. In the Section 5 we detail the implementation of branch-and-bound to be executed on a heterogeneous environment (CPUs and GPUs). In the Section 7 we report the experimental results, and finily, the consideration and future works in Section 8.

2. The Quadratic Assignment Problem (QAP) formulation

Given N facilities, N locations, a flow f_{ik} from each facility i to each facility k, $k \neq i$, and a distance d_{jn} from each location j to each location n, $n \neq j$, the QAP consists to assign each

facility i to exactly distinct location j in order to find:

$$min \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{\substack{k=1 \\ k \neq i}}^{N} \sum_{\substack{n=1 \\ n \neq j}}^{N} f_{ik} d_{jn} x_{ij} x_{kn} \tag{1}$$

$$\text{s.t.} \sum_{i=1}^{N} x_{ij} = 1 \ \ \forall \ \ \ j = 1, ..., N \tag{2}$$

$$\sum_{j=1}^{N} x_{ij} = 1 \ \ \forall \ \ i = 1, ..., N \tag{3}$$

$$x_{ij} \in \{0, 1\} \ \ \forall \ \ i = 1, ..., N; \ \ \ j = 1, ..., N \tag{4}$$

3. The Reformulation and linearization technique - RLT

For problems involving n variables, the technique of reformulation and linearization (RLT) provides n hierarchical levels of relaxation for a convex polygon of integer solutions to the original problem. For a given level k, $k \in \{1, ..., n\}$, also called as RLT k, the technique uses every polynomial factors of degree k involving k binary variables x or its complementaries $(1–x)$.

The linearization consists in adding auxiliary variables, each representing a possible product of original or complementary variables, and relating them to the original by new restrictions. Each new restriction corresponds to a unique restriction multiplied by a product of unique or additional variables. Assuming $c_{ijkn} = f_{ik} d_{jn}$ and $d_{ijknpq} = 0$ for all $i, k, p = 1, ..., N$ with distinct i, k, p, and $j, n, q = 1, ..., N$, also with distinct j, n, q. The following formulation is obtained when applying the method RLT2 to the QAP formulation (equations (1-4)), as Adams *et al*. (2007);

$$min \left\{ \sum_{i=1}^{N} \sum_{j=1}^{N} b_{ij} x_{ij} + \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{\substack{k=1 \\ k \neq i}}^{N} \sum_{\substack{n=1 \\ n \neq j}}^{N} c_{ijkn} x'_{ijkn} + \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{\substack{k=1 \\ k \neq i}}^{N} \sum_{\substack{n=1 \\ n \neq j}}^{N} \sum_{\substack{p=1 \\ p \neq i,k}}^{N} \sum_{\substack{q=1 \\ q \neq j,n}}^{N} d_{ijknpq} x''_{ijknpq} \right\} \tag{5}$$

$$\text{s.t.} \sum_{\substack{k=1 \\ k \neq i}}^{N} x'_{ijkn} = x_{ij} \ : \ (i, j, n = 1, .., N) \ ; \ n \neq j \tag{6}$$

$$\sum_{\substack{q=1 \\ j \neq n}}^{N} x'_{ijkn} = x_{ij} \ : \ (i, j, k = 1, .., N) \ ; \ i \neq k \tag{7}$$

$$\sum_{\substack{p=1 \\ p \neq i,k}}^{N} x''_{ijknpq} = x'_{ijkn} \ : \ (i, j, k, n, q = 1, .., N) \ ; \ k \neq i; \ \ distinct \ j, q, n \tag{8}$$

$$\sum_{\substack{q=1 \\ q \neq j,n}}^{N} x''_{ijknpq} = x'_{ijkn} \ : \ (i, j, k, n, p = 1, .., N) \ ; \ distinct \ i, k, p \ ; \ n \neq j \tag{9}$$

$$x'_{ijkn} = x'_{knij} \ (2 \, complementary \, coefficients) \ : \ (i, j, k, n = 1, .., N) \ ; \ i < k \ ; \ j \neq n \tag{10}$$

$$x''_{ijknpq} = x''_{ijpqkn} = x''_{knijpq} = x''_{knpqij} = x''_{pqijkn} = x''_{pqknij} \ (6\ complementary\ coefficients) : \\ (i,j,k,n,p,q = 1,..,N)\ ;\ i < k < p\ ;\ distinct\ j,n,q \tag{11}$$

$$x_{ij} \geq 0\ :\ (i,j = 1,..,N)\ ; \tag{12}$$

$$x'_{ijkn} \geq 0\ :\ (i,j,k,n = 1,..,N)\ ;\ i < k\ ;\ j \neq, n \tag{13}$$

$$x''_{ijknpq} \geq 0\ :\ (i,j,k,n,p,q = 1,..,N)\ ;\ i < k < p\ ;\ distinctj,n,q \tag{14}$$

and equations (2) and (3);

To obtain the formulation (5 - 14), we first apply the steps of RLT1: It's applied $(N \times (N-1))^2$ new constraints (6) and (7) by multiplying the constraints (2) and (3) for each of N^2 binary variables x_{kn}, , with $i \neq k$ and with $j \neq n$, (and changing the indexes i and k and the indexes j and n). After, we replaced each product $x_{ij}x_{kn}$ by binary variable x'_{ijkn} and each $x_{ij}x_{ij}$ for x_{ij}. Because the commutative property of the product, the additional constraints (10) are imposed. The x'_{ijkn} and x'_{knij} are called of complementary coefficients.

Next, we applied the steps of RLT2: It's added $(N \times (N-1) \times (N-2))^2$ news constraints (8) and (9) obtained by multiplying the constraints (6) and (7) by each of N^2 binary variables x_{pq}, with distinct i,k,p and also distinct j,n,q , (changing the indexes i, k and p, and the indexes j, n and q). Then, each product $x'_{ijkn}x_{pq}$ is replaced for each binary variable x''_{ijknpq}. Similarly the reformulation of RLT1, the constraint (11) also should be imposed. The coefficients $x''_{ijknpq}, x''_{ijpqkn}, x''_{knijpq}, x''_{knpqij}, x''_{pqijkn}$, and x''_{pqknij} are called complementary coefficients.

We represent the current dual solution by a set of matrices containing the modified coefficients costs (reduced costs) that are maintained non-negative during the execution of algorithm. The cost coefficients $b_{ij}\ \forall\ (i,j = 1,..,N)$ are stored in a $N \times N$ matrix B, the cost coefficients $c_{ijkn}\ \forall\ (i,j,k,n = 1,..,N)$, with $i \neq k$ and $j \neq n$ are stored in a $N(N-1) \times N(N-1)$ matrix C, and finally, the cost coefficients $d_{ijknpq}\ \forall\ (i,j,k,n,p,q = 1,..,N)$, with distinct i,k,p and also distinct j,n,q, are stored in a $N(N-1)(N-2) \times N(N-1)(N-2)$ matrix D.

4. The dual algorithm to calculate the lower bound

The Lower Bound, LB, is the resulting value from maximizing the objective function obtained from the dualization of the RLT2 formulation applied to the QAP (Equations (2), (3) and (5 - 14)). The dual algorithm implemented in this work consists of to modify the LB and the coefficients of B, C and D so that no value becomes negative and the cost of some viable solution to the QAP remains unchanged after modification according to the equations (2), (3) and (5 - 14). As result of this property, the LB, at any time of algorithm execution, is a valid bound for the cost of the optimal solution.

Our approach to maximize LB is to transfer costs from D to C, from C to B, and then, from B to LB. For this, we use a modified version of the dual RLT2 algorithm of Adams *et al.* (2007). The Algorithm 1 consists of a loop with 3 operations: Costs Spreading, Transfer between Complementary Costs and Cost Concentration.

The Parameter K is provided at the beginning of the execution of the application and serves to stop the dual loop. The K is the minimum percentage of the LB progress. The K doesn't have a defined value, it depends of the instance, dimension and the history of experiments. The K value is between 0,01% $(K = 0,00001)$ and $100\%(K = 1)$.

Algorithm 1:

begin
- $LB \leftarrow 0$
- $UB \leftarrow$ *best known solution researched by a heuristic*
- $K \leftarrow$ *minimal progress limit of* $LB(0,01\%ofUB)$
- $b_{ij} \leftarrow$ *(assignment cost of* i *at* $j \; \forall \; (i,j)$
- $c_{ijkn} \leftarrow f_{ik} \times d_{jn} \; \forall \; (i,j,k,n)$ *with* $i \neq k$ *and* $j \neq n$
- $d_{ijknpq} \leftarrow 0 \;\; \forall \; (i,j,k,n,p,q)$ *with distinct* i,k,p *and also distinct* j,n,q
- $progress \leftarrow 1$
- ***loop*** *While (*$progress >= K$*) and (*$LB < UB$*)*
  - *Costs Spreading from* B *to* C
  - *Costs Spreading from* C *to* D
  - *Transfer between Complementary Costs of* D
  - *Cost Concentration from* D *to* C *(*$c_{ijkn} \leftarrow Concentrate(d_{ijkn})$*)*
  - *Transfer between Complementary Costs of* C
  - *Cost Concentration from* C *to* B *(*$b_{ij} \leftarrow Concentrate(c_{ij})$*)*
  - *Cost Concentration from* B *to* LB' *(*$LB' \leftarrow Concentrate(B)$*)*
  - *LB* $\leftarrow$ *LB + LB'*
  - $progress \leftarrow (LB'/UB)$
- **End**

end

4.1. Cost Concentration

The cost concentration operation between the matrices consists of cost transfer from D to C, from C to B and from B matrix to LB, obeying the constraints (2 - 3) and (6 - 9). The cost concentration operation is solved as a linear assignment problem. To solve each linear assignment Adams *et al*. (2007) adopts the hungarian algorithm, Munkres (1957). The hungarian algorithm minimize the total cost S of a designation. We consider an example adopting a M cost matrix with dimension $N \times N$, where each coefficient M_{rs} corresponds to the cost of assigning a facility r to a location s. The objective function is then: $min \; S = \sum_r^N \sum_s^N M_{rs} x_{rs}$ where $x_{rs} \in \{0,1\}$, $\sum_{r=1}^N x_{rs} = 1 \; \forall \; s \in \{1,..,N\}$, $\sum_{s=1}^N x_{rs} = 1 \; \forall \; r \in \{1,..,N\}$.

For cost concentration operation from D to C, the M matrix has a dimension $(N-2) \times (N-2)$. For each (i,j,k,n), with $i \neq k$ and $j \neq n$, M receive the $(N-2)^2$ costs elements of D_{ijkn} submatrix. For each $(r,s = 1,..,N-2)$, M_{rs} receives $d_{(ijkn)pq}$ $(p,q = 1,..,N)$, with $p \neq \{i,k\}$ and $q \neq \{j,n\}$. S is obtained by the application of hungarian algorithm at M and adding S in c_{ijkn}. Finally, for each (i,j,k,n), with $i \neq k$ and $j \neq n$, each cost coefficient $d_{(ijkn)pq}$ is replaced by its correspondent residual coefficient of M. We represent this cost transfer operation as: $c_{ijkn} \leftarrow Concentrate(D_{ijkn})$. Similarly, it's done for cost concentration operation from C to B, represented by $b_{ij} \leftarrow Concentrate(C_{ij})$, with a $(N-1)^2$ M matrix and the cost concentration operation from B to LB, with a N^2 M matrix, represented this operation by $LB \leftarrow Concentrate(B)$.

4.2. Spreading costs

The cost spreading operation is the inverse of cost concentration operation, and is done from B to C and from C to D. The cost spreading operation from B to C consists of: For each (i,j), the cost element b_{ij} is spreading by the $(N-1)$ lines of C_{ij} submatrix, ie each cost element c_{ijkn} receives an increase of $b_{ij}/(N-1)$, $\forall \; (i,j,k,n)$, with $k \neq i$ and $n \neq j$. After the update of C, $b_{ij} = 0 \; \forall \; (i,j)$.

Similarly the previous operation, the cost spreading from C to D consists in: for each (i,j,k,n), the cost element c_{ijkn} is spreading by the $(N-2)$ lines of D_{ijkn} submatrix, ie each cost

element d_{ijknpq} receives an increase of $c_{ijkn}/(N-2)$, $\forall\ (i,j,k,n,p,q)$, with distinct i,k,p and also distinct j,n,q. After the update of D, $c_{ijkn} = 0\ \forall (i,j,k,n)$, with $k \neq i$ and $n \neq j$.

4.3. Transfer costs between complementary coefficients

The complementary coefficients enable communication costs between submatrices. This operation is essential to achieving a good LB by applying the RLT method. A good choice of cost transfer strategy allows to reach tighter LB with low runtime of dual algorithms. The charge transfer operation between complementary consists of increment one or more coefficients and decrement the other(s), maintaining the total of decrements and increments, between complementary coefficients, always equal to 0.

5. Considerations about the GPU

The GPU works as a co-processor for the CPU executing massively mathematical calculations. The GPUs are composed of several multiprocessor (or SM) type Single Instruction Multiple Data (SIMD). Each SM has a group of processing units (or SP). The NVidia C2070, for example, has 14 SMs each with 32 SPs.

The programming model Compute Unified Device Architecture (CUDA) enables the development of programs directed to exploit the potential of GPUs. The jobs are submitted by the CPU (host) for the GPU (device) through calls with signaling defined by the programming model. These tasks are called kernel. Each CUDA kernel is executed by a thread. These threads are grouped into blocks, and these in grids. When a CUDA program calls a grid to execute on the GPU, each blocks of this grid is directed to an available SM. Upon receiving a block, the SM divides the block in sets of 32 consecutive threads (warps). Each warp executes a single instruction at a time. When a block is finished, another is assigned to the SM.

The GPU memory is organized in a hierarchical way comprising of 4 distinct levels: the *Local* memory present in every SP, the *Shared* memory that can be accessed by any SP of the same SM, *Global* memory that is visible by any SP of all SMs, and finally, the *Constant* that only is accessible for reading by all SPs. Such levels have response times about 2, 2, 600 and 600 clock cycles, respectively.

6. Branch-and-bound with the dual RLT2 adapted for execution on GPU

The application containing the branch-and-bound algorithm is executed in the CPU (Host), creating a amount of cpu_threads (we use this denomination to trivial threads in order to differentiate of GPU threads) equal the total of available GPUs. The root node of the branch-and-bound tree is performed by cpu_thread with ID = 0. Next, at the first Branch, each cpu_thread takes a subtree (or node) and execute a depth-first search. When it finished your subtree, the cpu_thread takes another node that has not been fathomed.

The B, C and D matrices are stored in GPU Global memory. For each fathomed node the cpu_thread executes the dual algorithm of Section 4. The cost concentration, costs spreading and costs transfer between complementary coefficients operations are executed by GPU.

The cost concentration use a linear assignment algorithm which could be the Hungarian algorithm, Munkres (1957), widely used due to its efficiency, but this algorithm doesn't allow a good parallelization. We choose to use the Auction Algorithm Bertsekas e Castanon (1989), also used as a linear assignment algorithm and is more efficient for parallelization on GPU. Each person (auction participant) and each object of Auction algorithm corresponds to, here in this work, the facility and the location of QAP original problem, respectively. The parallelization strategy is to let each GPU thread with the function of a person on Auction Algorithm.

For the cost concentration operation from D to C, a $(N-2) \times (N-2)$ M matrix is allocated in Shared memory of each SM. For each (i,j,k,n), the cost coefficients of submatrix D_{ijkn} are transferred to M and a auction algorithm is executed. It's allocated one warp (32 threads) for each auction algorithm execution, and at the end, the residual coefficients of M matrix are transferred, from the Shared memory to D_{ijkn} on Global memory. This procedure is same for the cost concentration operation from C to B and from B to LB.

We opted not to implement the sharing of complementary coefficients at the same memory location, possible by the constraints (10) and (11). This resourse was used in Adams *et al*. (2007), and could reduce the required memory to store the D matrix. The reason of our option is the loss of application performance, since the coefficients of the matrices would not be in contiguous blocks, that would spend several read cycles to transfer from Global memory to M matrix in Shared memory. Another reason, which also causes loss of performance, is that several threads of distinct warps concurrently accessing the same location in memory, and thus, serializing access to memory.

We note that each coefficient of submatrix $(d_{ijkn})_{pq}$ corresponds to complementary coefficient of submatrix $(d_{knij})_{pq}$. Then, we introduced the concept of complementary submatrices where: for all $(i, k = i, ..., N)$, and $(j, n = 1, ..., N)$, with $i < k$, the submatrices $(d_{ijkn})_{pq}$ and $(d_{knij})_{pq}$ are complementaries. Using this concept, we can reduce by a half the required memory to store the D matrix, as also reduce by half the total of costs concentration operations from D to C impacting the application runtime.

The Branch strategy consists of a strong branch that is done a preliminary evaluation of a set of assignments and concentrations of the facilities not yet assigned to each location also unassigned, and then, it is selected a row or column that obtains the highest LB. In the strong branch the costs concentration follows the RLT1 dual algorithm similar to Section 4, except by the operations with transfer costs of D matrix.

7. Experimental Results

The proposed algorithm, here called GPU dual RLT2, was implemented using the C++ language and the CUDA programming model. The experiments were executed on a non-exclusive machine with a Intel Xeon Hexcore CPU with 24 GB of RAM and one NVidia GeForce GTX TITAN GPU. This GPU has 3,072 of 1,000 MHz SPs, organized in 24 SMs (128 SPs each) and 12 GB DDR5 Global memory.

Table 1 shows the results achieved for different QAP instances. We used 2 others works as reference to evaluate our experiments, the sequential dual RLT2 algorithm of Adams *et al*. (2007) and the RLT1/2/3 Parallel C algorithm of Hahn *et al*. (2013). Recent experiments were done with the dual RLT2 algorithm and the new results were published in Hahn *et al*. (2012) and we put them in Table 1. The dual RLT2 algorithm experiments of Hahn *et al*. (2013) were made in two environments: one machine with a 1.9GHz E6900 CPU, used to process the Nug20, Nug22, Nug24 and Nug25 instances, and one machine with a 733MHz Itanium CPU used to process the Nug28 and Nug30 instances. The RLT1/2/3 Parallel C algorithm experiments of Hahn *et al*. (2013) were executed in 30 of the 64 hosts of Palmetto Supercomputing Cluster at the University of Clemson, with 2TB of shared memory and each host with one 2.2GHz Intel Xeon CPU.

Our proposal was evaluated with instances found in QAPLIB, Burkard *et al*. (1997). Besides the instances used in Hahn *et al*. (2013) we added others as: Tai20a, Tai20b, Tai25a, Tai25b, Tai30b, Tai35b, Tai40b, Kra30a, Kra30b and Tho30.

Table 1 is organized as follows: the initial three columns has the information about the instances as identification, the total of facilities (N) and the optimal value found in the literature. The other columns report the results of the algorithms cited above. In the fourth, fifth and sixth column, the amount of nodes of branch-and-bound, the total of main memory and the runtime, respectively, for the dual RLT2 algorithm. In the seventh, eighth, ninth and tenth column, the number of nodes of branch-and-bound, the total of nodes in RLT3 step, the amount of main memory and the runtime, respectively. Finally, the results obtained in our algorithm (GPU dual RLT2), in sequence: the total of fathomed nodes, the amount of RAM used in the *host*, the total of Global memory used in GPU, and the runtime clock.

Comparing the results achieved by dual RLT2 and the GPU dual RLT2 we observed that the speedup is between 50 to 300 depending of the instance. Some instances have the B&B subtrees very unbalanced that some subtrees are finished earlier than others, then, load balancing procedures were implemented.

Instance	N	Optimal	dual RLT2			RLT1/2/3 Parallel C				GPU dual RLT2			
			Nodes	Memory	time	Nodes		Memory	time	Nodes	Mem. (GB)		time
			B&B	(GB)	(s)	B&B	RLT3	(GB)	(s)	B&B	*Host*	*Device*	(s)
nug20	20	2570	1,407	0.4	2,978	39	9	18.4	736	565	1	0.4	6
nug22	22	3596	1,450	0.8	3,361	52	16	40.7	1,578	927	1	0.6	11
nug24	24	3488	37,099	0.8	5,781	102	16	85.0	5,097	606	1,5	1.2	29
nug25	25	3744	15,497	1.7	124,702	267	31	120.3	23,072	3,616	2	1.5	491
nug27	27	5234	-	-	-	359	46	231.5	53,227	1,831	3	2	360
nug28	28	5166	202,295	3.6	2,856,392	1,538	115	321.4	130,527	7,269	4	3	2,084
nug30	30	6124	543,061	6	19,735,563	3,383	251	722.8	733,812	103,380	5	6	88.867
tai20a	20	703482	-	-		-	-	-	-	7,117	1	0.4	71
tai20b	20	122455319	-	-		-	-	-	-	211	1	0.4	6
tai25a	25	1167256	-	-	-	-	-	-	-	324,394	2	1.7	109,011
tai25b	25	344355646	-	-	-	-	-	-	-	350	2	1.7	63
tai30b	30	637117113	-	-	-	W/O inf	W/O inf	753	315.584	1,011	5	6	448
tai35b	**35**	**283315445***	-	-	-					**1,029,312**	**5**	**9**	**473,873**
tai40b	**40**	**637250948***	-	-	-					**2,510,362**	**5**	**12**	**4,949,444**
kra30a	30	88900	-	-	-	-	-	-	-	20,764	4	6	940
kra30b	30	91420	-	-	-	-	-	-	-	24,082	4	6	16,552
tho30	30	149936	-	-	-	-	-	-	-	22,429	4	6	17,219

*** Instances solved exactly for the first time**

Table 1: Comparison between the results of Hahn *et al*. (2012) with the dual RLT2 of Adams *et al*. (2007), the results of RLT 1/2/3 Parallel C algorithm of Hahn *et al*. (2013) and the results of GPU dual RLT2 algorithm proposed in this work

The RLT1/2/3 Parallel C algorithm extends the relaxation to the level 3 RLT reaching tighter bounds that dual RLT2, however, there is the need of an extensive working memory. Besides the B, C and D matrices of RLT2, the implementation of RLT3 requires a $(N \times (N-1) \times (N-2) \times (N-3))^2$ E matrix that to require a large amount of working memory for storing and processing the operations with this matrix .

Comparing the results of the GPU dual RLT2 algorithm with the achieved by the RLT1/2/3 Parallel C algorithm, we observed that the runtime is about 140x lower for instances as Nug20, Nug22, Nug24 and Nug27, approximately 50x lower for Nug25 and Nug28 and about of 8.5x lower for Nug30. The reduction in performance on experiments of RLT2 in relation to RLT3 with increasing size of the problem is pointed out in Zhu *et al*. (2012). They show that the runtime of one instance larger than 28 facilities in applications based in RLT2 relaxation are superior than applications with relaxation RLT3.

Instance	Optimal	Solution
tai35b	**283315445**	14 12 5 18 10 30 11 22 1 19 9 20 32 17 2 33 3 8 13 27 16 4 34 7 23 24 6 35 31 28 15 21 29 26 25
tai40b	**637250948**	36 1 15 11 25 37 31 19 39 13 27 7 40 22 4 33 16 34 10 14 12 23 5 32 35 38 9 3 30 29 24 17 2 6 28 8 20 26 18 21

Table 2: Instances solved exactly by first time

In Table 1, the Instances Tai35b and Tai40b in bold were resolved exactly by first time confirming the optimal value obtained by heuristics. The solutions reached are present in Table 2. To solve these instances the GPU dual RLT2 algorithm was executed on a machine without exclusivity and subject to interruptions due to lack of energy and decreased network, imposing the need to implement a checkpoints procedure. The runtime shown in Table 1 corresponds to the sum of the individual clock times of execution and each re-execution after the interrupts. The gap reached in Branch-and-bound root node also were the lower of literature, they were 4.92 % and 4.14 % for tai35b and tai40b, respectively, the best known found in QAPLIB site were 14.52 % and 11.43 %.

8. Conclusions and future works

The recent commercial GPUs have the tendence a high increasing of amount of SPs, however, their global memory have not the same growth, this scenario provides advantages to applications with relaxation RLT2 over the applications with RLT3.

The GPU dual RLT2 algorithm is still under review, and new changes now allow better performance. A distributed version is being evaluated promising good results. Results and new contributions will be presented in future work.

*Bibliography

Adams, W. P., Guignard, M., Hahn, P. M. e Hightower, W. L. (2007), A level-2 reformulation-linearization technique bound for the quadratic assignment problem. *European Journal of Operational Research*, v. 180, n. 3, p. 983–996.

Adams, W. P. e Sherali, H. D. (1990), Linearization strategies for a class of zero-one mixed integer programming problems. *Operations Research*, v. 38, n. 2, p. 217–226.

Bertsekas, D. e Castanon, D. (1989), The auction algorithm for the transportation problem. *Annals of Operations Research*, v. 20, n. 1, p. 67–96.

Burer, S. e Vandenbussche, D. (2006), Solving lift-and-project relaxations of binary integer programs. *SIAM Journal on Optimization*, v. 16, p. 726–750.

Burkard, R. E., Karisch, S. E. e Rendl, F. (1997), Qaplib - a quadratic assignment problem library. *J. of Global Optimization*, v. 10, n. 4, p. 391–403.

Burkard, R. E., Çela, E., Pardalos, P. M. e Pitsoulis, L. S. The quadratic assignment problem, 1998.

Cela, E. *The Quadratic Assignment Problem: Theory and Algorithms*. Combinatorial Optimization. Springer. ISBN 9780792348788, 1997.

Gonçalves, A., Drummond, L., Pessoa, A. e Hahn, P. (2013), Improving lower bounds for the quadratic assignment problem by applying a distributed dual ascent algorithm. v. .

Hahn, P., Roth, A., Saltzman, M. e Guignard, M. (2013), Memory-aware parallelized rlt3 for solving quadratic assignment problems.

Hahn, P. M. e Grant, T. (1998), Lower bounds for the quadratic assignment problem based upon a dual formulation. *Operations Research*, v. 46, n. 6, p. 912–922.

Hahn, P. M., Zhu, Y.-R., Guignard, M., Hightower, W. L. e Saltzman, M. J. (2012), A level-3 reformulation-linearization technique-based bound for the quadratic assignment problem. *INFORMS J. on Computing*, v. 24, n. 2, p. 202–209.

Koopmans, T. C. e Beckmann, M. (1957), Assignment problems and the location of economic activities. *Econometrica*, v. 25, n. 1, p. 53–76.

Munkres, J. (1957), Algorithms for the assignment and transportation problems. *Journal of the Society of Industrial and Applied Mathematics*, v. 5, n. 1, p. 32–38.

Padberg, W. e Rijal, P. *Location, Scheduling, Design and Integer Programming*. Kluwer Academic Publishers. ISBN 978-1-4612-8596-0, 1996.

Pardalos, P. M., Rendl, F. e Wolkowicz, H. The quadratic assignment problem: A survey and recent developments. *In Proceedings of the DIMACS Workshop on Quadratic Assignment Problems, volume 16 of DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, p. 1–42. American Mathematical Society, 1994.

Sherali, H. e Adams, W. Reformulation-linearization techniques for discrete optimization problems. Du, D.-Z. e Pardalos, P. (Eds.), *Handbook of Combinatorial Optimization*, p. 479–532. Springer US, 1999.

Tsutsui, S. e Fujimoto, N. Solving quadratic assignment problems by genetic algorithms with gpu computation: A case study. *Proceedings of the 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference: Late Breaking Papers*, GECCO '09, p. 2523–2530, New York, NY, USA. ACM. ISBN 978-1-60558-505-5, 2009.

Tsutsui, S. e Fujimoto, N. Aco with tabu search on a gpu for solving qaps using move-cost adjusted thread assignment. *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*, GECCO '11, p. 1547–1554, New York, NY, USA. ACM. ISBN 978-1-4503-0557-0, 2011.

Zhu, Y.-R., Hahn, P. M., Guignard, M., Hightower, W. L. e Saltzman, M. J. (2012), A level-3 reformulation-linearization technique-based bound for the quadratic assignment problem. *INFORMS J. on Computing*, v. 24, n. 2, p. 202–209.